

AN OUTLINE OF THE BAYESIAN DECISION THEORY

H.R.N. VAN ERP, R.O. LINGER, AND P.H.A.J.M. VAN GELDER

ABSTRACT. In this fact sheet we give an outline on the Bayesian Decision Theory.

1. INTRODUCTION

The Bayesian decision theory is very simple in structure. Its algorithmic steps are the following:

- (1) Use the product and sum rules of Bayesian probability theory to construct outcome probability distributions.
- (2) If our outcomes are monetary in nature, then by way of the Bernoulli utility function we may map utilities to the monetary outcomes of our outcome probability distributions.
- (3) Maximize the position of the resulting utility probability distributions.

This is the whole of the Bayesian decision theory.

2. CONSTRUCTING OUTCOME PROBABILITY DISTRIBUTIONS

In the Bayesian decision theory each problem of choice is understood to consist of a set of decisions from which we must choose. Each possible decision has associated with it its own set of possible outcomes, and each outcome has its own plausibility of occurring relative to the other outcomes under that same decision. Stated differently, each decision in our problem of choice admits its own outcome probability distribution.

In its most abstract form, we have that each problem of choice consists of a set of potential decisions

$$D_i = \{D_1, \dots, D_n\}.$$

Each decision D_i we make may give rise to a set of possible events

$$E_{j_i} = \{E_{1_i}, \dots, E_{m_i}\}.$$

These events E_{j_i} are associated with the decisions D_i by way of the conditional probabilities $P(E_{j_i} | D_i)$. Furthermore, each event E_{j_i} allows for a set of potential outcomes

$$O_{k_{j_i}} = \{O_{1_{j_i}}, \dots, O_{l_{j_i}}\}.$$

These outcomes $O_{k_{j_i}}$ are associated with the events E_{j_i} by way of the conditional probabilities $P(O_{k_{j_i}} | E_{j_i})$.

By way of the product rule [5], we compute the bivariate probability distribution of an event and an outcome conditional on the decision taken:

$$P(E_{j_i}, O_{k_{j_i}} | D_i) = P(E_{j_i} | D_i) P(O_{k_{j_i}} | E_{j_i}). \quad (2.1)$$

The outcome probability distribution is then obtained by marginalizing, by way of the sum rule [5], over all the possible events:

$$P(O_{k_{j_i}} | D_i) = \sum_{j_i=1}^{m_i} P(E_{j_i}, O_{k_{j_i}} | D_i). \quad (2.2)$$

The outcome probability distributions (2.2), for $i = 1, \dots, n$, are the information carriers which represent our state of knowledge in regards to the objective consequences of our decisions.

3. A CONSISTENCY PROOF OF THE BERNOULLI UTILITY FUNCTION

We will now derive the Bernoulli utility function, or, equivalently, the Weber-Fechner law, or, equivalently, in content, Steven's Power law, using the desiderata of invariance and consistency. In this we follow a venerable Bayesian tradition, [2, 5, 7].

Say, we have the positive quantities x , y , and z , of some stimulus or commodity of interest. Then these quantities, being numbers on the positive real, admit an ordering. So, let quantities be ordered as $x \leq y \leq z$. We now want to find the function f that quantifies the perceived decrease associated with going from, say, the quantity z to the quantity x .

The first functional equation is based on the desideratum that the unknown function f should be invariant for a change of scale in our quantities:

$$f(x, z) = f(cx, cz), \quad (3.1)$$

where c is positive constant.

For example, if our quantities concern sums of money, then the perceived loss of going from ten dollars to one dollar should be the same perceived loss if we reformulate this scenario in dollar cents.

The second functional equation is based on the desideratum of consistency, in which we state that the perceived decrease in going directly from z to x , ought to be the same perceived decrease in going from z to x via y :

$$f(x, z) = g[f(x, y), f(y, z)]. \quad (3.2)$$

For example, if our quantities concern sums of money, then the perceived loss of going from ten dollars to one dollar should be the same perceived loss if we first go from ten dollars to five dollars, and then from five dollars to one dollar; seeing that in both scenarios we start out with an initial wealth of ten dollars, only to end up with a current wealth of one dollar.

The general solution to (3.1) is [3]:

$$f(x, y) = h\left(\frac{x}{y}\right), \quad (3.3)$$

where h is some arbitrary function. The general solution to (3.2) is [7]:

$$\Theta[f(x, z)] = \Theta[f(x, y)] + \Theta[f(y, z)], \quad (3.4)$$

where Θ is some arbitrary monotonic function. Moreover, because of this arbitrariness, we may define Θ as [7]:

$$\Theta(x) = \log \Psi(x). \quad (3.5)$$

Using (3.5), we may rewrite (3.4), without any loss of generality, as

$$\log \Psi[f(x, z)] = \log \Psi[f(x, y)] + \log \Psi[f(y, z)], \quad (3.6)$$

or, equivalently,

$$\Psi[f(x, z)] = \Psi[f(x, y)] \Psi[f(y, z)]. \quad (3.7)$$

Substituting (3.3) into (3.4) and (3.7) and letting, respectively,

$$\theta\left(\frac{x}{y}\right) = \Theta\left[h\left(\frac{x}{y}\right)\right], \quad (3.8)$$

and

$$\psi\left(\frac{x}{y}\right) = \Psi\left[h\left(\frac{x}{y}\right)\right], \quad (3.9)$$

we obtain the equivalent functional equations:

$$\theta\left(\frac{x}{z}\right) = \theta\left(\frac{x}{y}\right) + \theta\left(\frac{y}{z}\right) \quad (3.10)$$

and

$$\psi\left(\frac{x}{z}\right) = \psi\left(\frac{x}{y}\right) \psi\left(\frac{y}{z}\right). \quad (3.11)$$

If we assume differentiability, then (3.10), together with the two boundary conditions:

$$f(x, x) = \theta\left(\frac{x}{x}\right) = 0, \quad (3.12)$$

and

$$f(x, y) = \theta\left(\frac{x}{y}\right) < 0, \quad \text{for } x < y, \quad (3.13)$$

is sufficient to find the function f that quantifies the perceived decrease associated with going from the quantity y to the quantity x .

This function θ turns out to be Bernoulli's utility function, or, equivalently, the Weber-Fechner law of sense perception:

$$f(x, y) = q \log \frac{x}{y}, \quad q \geq 0 \quad (3.14)$$

where y is our initial asset position and x is the final asset position, and q is some arbitrary constant which has to be obtained by way psychological experimentation.

So, Bernoulli's utility function (3.14) is the only function that adheres to the desiderata of unit invariance and consistency, respectively, (3.1) and (3.2), and the boundary conditions that a zero change should lead to a zero perceived loss and that a perceived loss should be assigned a negative value, respectively, (3.12) and (3.13). Any other utility function will be in violation with these fundamental desiderata and specific boundary conditions.

Note that Fechner re-derived (3.14) in 1860 as the law that guides our sensory perception. In the years that followed (3.14) proved to be so successful, as it, amongst other things, gave rise to our decibel scale, that it established psychology as a legitimate experimental science [4]. But as Fechner was very careful, for metaphysical reasons, or so we hazard to guess [3], to apply his Weber law, which later became the Fechner-Weber law, only to non-monetary stimuli, the implied universality of (3.14) was not recognized for the longest time.

However, because of the here given consistency derivation of (3.14), it is now shown that the Fechner-Weber, or, equivalently, Bernoulli's utility function, is one of the consistent functions that quantifies the distance between x and y ; thus, explaining the universal applicability of Bernoulli's utility function.

The other consistent distance function is Steven's power law, which may be derived as follows. If we assume differentiability, then (3.11), together with the two boundary conditions:

$$f(x, x) = \psi\left(\frac{x}{x}\right) = 1, \quad (3.15)$$

and

$$0 < f(x, y) = \psi\left(\frac{x}{y}\right) < 1, \quad \text{for } x < y, \quad (3.16)$$

is sufficient to find the function f that quantifies the perceived decrease associated with going from the quantity y to the quantity x .

This function f turns out to be Steven's power law:

$$f(x, y) = \left(\frac{x}{y}\right)^q, \quad q \geq 0. \quad (3.17)$$

where y is our initial asset position and x is the final asset position, and q is some arbitrary constant which has to be obtained by way psychological experimentation.

So, Steven's power law (3.17) is the only function that adheres to the desiderata of unit invariance and consistency, respectively, (3.1) and (3.2), and the boundary

conditions that a zero change should lead to a ratio of one between the initial and final ‘asset position’ and that a perceived loss should be assigned a value smaller than 1, respectively, (3.15) and (3.16). Any other utility function will be in violation with these fundamental desiderata and specific boundary conditions.

We summarize, given the desiderata (3.1) and (3.2), the Fechner-Weber law (3.14) results from the boundary condition that negative increments result negative utilities and a zero increment results in an utility of zero, (3.12) and (3.13); whereas Steven’s power law (3.17) results from the boundary condition that utilities must be greater than zero and that a zero increment results an utility of one, (3.15) and (3.16).

Stated differently, the Fechner-Weber law and Steven’s power law are both equivalent in content, differing only in the proposed utility scale. A subtlety that seems to have been overlooked by some, seeing that the Fechner-Weber law versus the Steven’s power law has been a source of controversy in psycho-physical community[8].

4. THE CRITERION OF CHOICE AS A DEGREE OF FREEDOM

Let D_1 and D_2 be two decisions we have to choose from. Let o_i , for $i = 1, \dots, n$, and o_j , for $j = 1, \dots, m$, be the monetary outcomes associated with, respectively, decisions D_1 and D_2 . Then in the Bayesian decision theory we first construct the two outcome distributions that correspond with these decisions:

$$p(o_i | D_1), \quad p(o_j | D_2), \quad (4.1)$$

where, if $n = m$, the outcomes o_i and o_j may or not may be equal for $i = j$.

We then proceed, by way of the Bernoulli utility function (3.14), or, equivalently, the Weber-Fechner law, to map utilities to the monetary outcomes o_i and o_j in (4.1). This leaves us with the utility probability distributions:

$$p(u_i | D_1), \quad p(u_j | D_2). \quad (4.2)$$

Now, our most primitive intuition regarding the utility probability distributions (4.2) is that the decision which corresponds with the utility probability distribution which lies more to the right will also be the decision that promises to be the most advantageous. So, when making a decision we ought to compare the positions of the utility probability distributions on the utility axis and then choose that decision which maximizes the position of these utility probability distributions.

This all sounds intuitive enough. But how do we define the position of a probability distribution? Ideally we would have some consistency derivation of what constitutes a position measure of a probability distribution, say,

$$H_n(p_1, \dots, p_n, x_1, \dots, x_n) \quad (4.3)$$

where p_i are the probabilities of the values x_i , for $i = 1, \dots, n$. But in the absence of such a consistency derivation we have to take our recourse to *ad hoc* common sense considerations. Stated differently, the criterion of choice in our decision theory constitutes a degree of freedom.

4.1. The Expectation Value as a Position Measure. From the introduction of expected outcome theory in the 17th century and expected utility theory in the 18th century the implicit assumption has been that the measure of a position of a probability distribution is given by its expectation value [5, 1]:

$$E(X) = \sum_{i=1}^n p_i x_i = H_n(p_1, \dots, p_n, x_1, \dots, x_n) \quad (4.4)$$

But this criterion of choice has proven to be so unsatisfactory that it has given rise to the paradigm of behavioral economics which holds as its central tenet that human decision making does not adhere to the maximization of expectation values [6]. So, we set out to search for a more appropriate criterion of choice.

4.2. The Confidence Bounds as a Position Measure. Now we may imagine a decision problem in which we are only interested in the positions of the probabilistic worst or best case scenarios.

The absolute worst case scenario is:

$$a = \min(x_1, \dots, x_n). \quad (4.5)$$

The criterion of choice (4.5) is also known as the maximin criterion of choice.

The k-sigma lower bound of a given probability distribution is given as

$$LB(X) = E(X) - k \text{std}(X), \quad (4.6)$$

where

$$\text{std}(X) = \sqrt{\sum_{i=1}^n p_i [x_i - E(X)]^2}, \quad (4.7)$$

and where k is the sigma level of the lower bound. The probabilistic worst case scenario then is given as:

$$LB^*(X) = \begin{cases} E(X) - k \text{std}(X), & LB(X) > a, \\ a, & LB(X) \leq a. \end{cases} \quad (4.8)$$

So, we have that the probabilistic worst case scenario holds the maximin criterion of choice as a special case.

The absolute best case scenario is:

$$b = \max(x_1, \dots, x_n). \quad (4.9)$$

The criterion of choice (4.9) is also known as the maximax criterion of choice.

The k -sigma upper bound of a given probability distribution is given as:

$$UB(X) = E(X) + k \text{std}(X), \quad (4.10)$$

where k is the sigma level of the upper bound. The probabilistic best case scenario then is given as:

$$UB^*(X) = \begin{cases} E(X) + k \text{std}(X), & UB(X) < b, \\ b, & UB(X) \geq b. \end{cases} \quad (4.11)$$

So, we have that the probabilistic best case scenario holds the maximax criterion of choice as a special case.

If we take as our criterion of choice (4.8) then we only endeavor to minimize our ‘losses’ and if we take as our criterion of choice (4.11) then we only endeavor to maximize our ‘gains’. A more rational, that is, balanced, criterion of choice would be to make a trade-off between the losses/gains in the probabilistic worst case scenarios (4.8) and the corresponding gains/losses in the probabilistic best case scenarios (4.11).

4.3. The Sum of Confidence Bounds as a Position Measure. If we take as our criterion of choice

$$\frac{LB^*(X) + UB^*(X)}{2} = \begin{cases} E(X), & LB(X) > a, UB(X) < b, \\ \frac{a + E(X) + k \text{std}(X)}{2}, & LB(X) \leq a, UB(X) < b, \\ \frac{E(X) - k \text{std}(X) + b}{2}, & LB(X) > a, UB(X) \geq b, \\ \frac{a + b}{2}, & LB(X) \leq a, UB(X) \geq b, \end{cases} \quad (4.12)$$

then we have a position measure which makes a trade-off between the losses/gains in the probabilistic worst case scenarios (4.8) and the corresponding gains/losses in the probabilistic best case scenarios (4.11); see Appendix A.

This alternative position measure, as an added benefit, also holds the traditional criterion of choice (4.4) as a special case, when no undershoot and overshoot of, respectively, the lower and upper sigma confidence bounds occur, as well as Hurwicz’s criterion of choice with a balanced pessimism factor of $c = 1/2$, when both an undershoot and an overshoot occur. Nonetheless, it may be found that the criterion of choice (4.12) is vulnerable to a simple counter-example.

Imagine two utility probability distributions having equal lower and upper bounds, but one being right-skewed and the other being left-skewed. Then the criterion of choice (4.12) will leave us undecided between the two, whereas our intuition would give preference to the decision corresponding with the left-skewed distribution, as the bulk of the probability distribution of the left-skewed distribution will be more to the right than that of the right-skewed distribution.

4.4. The Sum of Confidence Bounds Plus the Expectation Value as a Position Measure. What we seek to maximize in our decision theory is the position of the utility probability distributions; as we have that the decision that puts our utility probability distribution most to the right promises to be the most profitable decision. In this there is little room for maneuvering. But in our choice of the measure that captures the position of a given probability distribution there is all the more.

Taking a cue from the behavioral economists we have derived as an alternative to (4.4) the criterion of choice (4.12) that also takes into account the standard deviation of a given probability distributions, by way of the positions of the under and overshoot corrected lower and upper bounds. But only to find its universality compromised by the simple counter example of a right-skewed and a left-skewed distribution which have the same lower and upper bounds.

Now, also taking a cue from the intuitive results which flow forth from (4.12) [3], we may ‘repair’ our criterion of choice (4.12), albeit in an *ad hoc* fashion, by taking as our position measure for a probability distribution the weighted sum:

$$\frac{LB^*(u) + E(u) + UB^*(u)}{3} = \begin{cases} E(X), & LB(X) > a, UB(X) < b, \\ \frac{a+2E(X)+k \text{ std}(X)}{3}, & LB(X) \leq a, UB(X) < b, \\ \frac{2E(X)-k \text{ std}(X)+b}{3}, & LB(X) > a, UB(X) \geq b, \\ \frac{a+E(X)+b}{3}, & LB(X) \leq a, UB(X) \geq b, \end{cases} \quad (4.13)$$

For in this criterion of choice we not only take into account the trade-off between the probabilistic worst and best case scenarios, but also the location of the bulk of the probability density in a uni-model probability distribution; thus, accommodating the intuitive preference for the left-skewed distribution of the above counter example.

The position measure (4.13) is the weighted sum of the positions of, respectively, the probabilistic worst, expected, and best case. The uncorrected lower and upper bounds, (4.6) and (4.10), have been traditionally used as simplifying proxies for their generating probability distributions, by way of confidence intervals:

$$[LB(X), UB(X)]. \quad (4.14)$$

We, alternatively, take as our simplifying proxy the corrected lower and upper bounds, (4.8) and (4.11), and the expectation value (4.4):

$$[LB^*(X), E(X), UB^*(X)]. \quad (4.15)$$

Because of the corrections for lower bound undershoot and upper bound overshoot in (4.15) we have that for skewed distributions the distance between $E(X)$ and

$LB^*(X)$ may differ from the distance between $UB^*(X)$ and $E(X)$; thus, reflecting the asymmetry present in these distributions.

The position of the generating probability distribution then is taken to be the weighted sum of the positions of the elements of our simple proxy distribution. This then is the rationale behind the criterion of choice (4.13).

5. DISCUSSION

It may be read in Jaynes' [5], that to the best of his knowledge, there are as of yet no formal principles at all for assigning numerical values to loss functions; not even when the criterion is purely economic, because the utility of money remains ill-defined. In the absence of these formal principles, Jaynes final verdict was that decision theory can not be fundamental.

The Bernoulli utility function, initially derived by Bernoulli, by way of common sense first principles [1], has now been derived by way of a consistency argument. This consistency argument explains why it is that Bernoulli's utility function, both in its original Fechner-Weber law and in its alternative Steven's power law form, has proven to be so ubiquitous and successful the field of sensory perception research; simply because human sense perception, like the laws of Nature, adheres to the desideratum of consistency.

The history of Bayesian probability theory has taught us that the usefulness of a theory, in terms of its practical and beautifully intuitive results, in the absence of a compelling axiomatic basis, provides no safeguard against attacks by those who choose to close their eyes to this usefulness. This is why we felt compelled to search for a consistency derivation of the Bernoulli utility function.

Now, having presented a consistency proof for the Bernoulli utility function, the question now is: Is the Bayesian decision theory, just like the Bayesian probability and information theories, Bayesian in the strictest sense in the word, or, equivalently, an inescapable consequence of the desideratum of consistency? We will now try to answer this question.

The first two algorithmic steps of the Bayesian decision theory, respectively, the construction of outcome probability distributions by way of the Bayesian probability theory and the construction of utility probability distributions by way of the Bernoulli utility function, allow us no freedom.

To construct our outcome and utility probability distributions otherwise, would be to invite inconsistency. But there is one degree of freedom remaining in the Bayesian decision theory as a whole. This remaining degree of freedom lies in the choice of our position measure of a given probability distribution.

In any problem of choice we will endeavor to choose that decision which has a corresponding utility probability distribution that is lying most the right on the

utility axis; that is, we will choose to maximize our utility probability distributions. In this there is little freedom. But we are free, in principle, to choose the measures of the positions of our utility probability distributions any way we see fit. Nonetheless, we believe that it is always a good policy to take into account all the pertinent information we have.

If we only maximize the expectation values of the utility probability distributions, then we will, by definition, neglect the information that the standard deviations of the utility probability distributions have to bear on our problem of choice, by way of the symmetry breaking in the case of an overshoot of one of the bounds.

Likewise, we are free to only maximize one of the confidence bounds of our utility probability distributions, while neglecting the other. But in doing so, we will be performing probabilistic maximin or maximax analyses, and, consequently, neglect the possibility of either the (catastrophic) losses in the lower bound or the (astronomical) gains in the upper bound.

However, if we only maximize the sum of the lower and upper bound, or a scalar multiple thereof, then we will make a trade-off between the probabilistic worst and best case scenarios. But in the process, we will, for uni-modal distributions, be neglecting the location of the bulk of our probability distributions.

This is why, in our minds, the scalar multiple the sum of the lower bound, expectation value, and upper bound currently is the most all-round position measure for a given probability distribution, as it reflects the position of the probabilistic worst and best case scenarios, as well as the position of the expected outcome.

Having removed the degree of freedom of the utility function by way of a consistency derivation, we now should endeavor to find a like consistency derivation of the measure of the position of a given probability distribution (4.3); as such a consistency derivation would make the Bayesian decision theory incontestable.

But until that time, we will have to do with the kind of simplistic common sense reasoning that led us from the traditional position measure (4.4), to the position measures (4.12) and (4.13), and make the disclaimer that our adopted criterion of choice, as a degree of freedom, is just a matter of choice.

Acknowledgments: We would like here to express our gratitude to Kevin H. Knuth, whose kind and patient feedback led us to our consistency proof of Bernoulli's utility function and to Kevin M. Vanslette, whose simple but highly effective 'counter' example led us to drop (4.12) and propose (4.13) instead. The research leading to these results has received partial funding from the European Commission's Seventh Framework Program [FP7/2007-2013] under grant agreement no.265138.

REFERENCES

- [1] Bernoulli D.: *Exposition of a New Theory on the Measurement of Risk*. Translated from Latin into English by Dr Louise Sommer from ‘Specimen Theoriae Novae de Mensura Sortis’, *Commentarii Academiae Scientiarum Imperialis Petropolitanae*, Tomus V, 175-192, (1738).
- [2] Cox R.T.: *Probability, Frequency and Reasonable Expectation*, *American Journal of Physics*, 14, 1-13, (1946).
- [3] Erp van H.R.N., Linger R.O., and Gelder van P.H.A.J.M.: *Fact Sheet on the Bayesian Decision Theory*, arXiv, (2015).
- [4] Fancher R.E.: *Pioneers of Psychology*, W. W. Norton and Company, London, (1990).
- [5] Jaynes E.T.: *Probability Theory; the Logic of Science*. Cambridge University Press, (2003).
- [6] Kahneman D.: *Thinking, Fast and Slow*, Penguin Random House, UK, (2011).
- [7] Knuth K.H. and Skilling J.: *Foundations of Inference*, arXiv: 1008.4831v1 [math.PR], (2010).
- [8] Stevens S.S.: *To Honor Fechner and Repeal His Law*, *Science*, New Series, Vol. 133, No. 3446, pp 80-86, (1961).

APPENDIX A. DERIVING THE SUM OF THE LOWER AND UPPER CONFIDENCE BOUND MEASURE

Now, the confidence bounds of (4.2), say:

$$[LB(u|D_1), UB(u|D_1)], \quad [LB(u|D_2), UB(u|D_2)], \quad (A.1)$$

may provide us with a numerical handle on the concept of more-to-the-right.

For example, if we have that both

$$LB(u|D_1) > LB(u|D_2), \quad UB(u|D_1) > UB(u|D_2). \quad (A.2)$$

Then we will have an unambiguous preference for decision D_1 over decision D_2 ; seeing that under both the still probable worst and best case we will be better if we opt for D_1 .

Likewise, if we have that either

$$LB(u|D_1) = LB(u|D_2), \quad UB(u|D_1) > UB(u|D_2), \quad (A.3)$$

or

$$LB(u|D_1) > LB(u|D_2), \quad UB(u|D_1) = UB(u|D_2). \quad (A.4)$$

Then, again, we will have an unambiguous preference for decision D_1 over decision D_2 . In the constellation (A.3), we stand, all other things being equal, to be better of under the still probable best case scenario; while in the constellation (A.4), we stand, all other things being equal, to be less worse of under the still probable worst case scenario.

However, things become more ambiguous when, say, under decision D_1 , we have to make a trade-off between either a gain in the upper bound and a loss in the lower bound

$$LB(u|D_1) < LB(u|D_2), \quad UB(u|D_1) > UB(u|D_2), \quad (A.5)$$

or a gain in the lower bound and a loss in the upper bound

$$LB(u|D_1) > LB(u|D_2), \quad UB(u|D_1) < UB(u|D_2). \quad (\text{A.6})$$

We *postulate* here that a rational criterion of choice in the respective trade-off situations (A.5) and (A.6), would be to pick that decision whose gain in either the lower or upper bound exceeds the loss in the corresponding upper or lower bound.

So, if, say, under D_1 we stand to gain more in the still probable best case scenario than we stand to lose under the still probable worst case scenario, that is, (A.5):

$$LB(u|D_2) - LB(u|D_1) < UB(u|D_1) - UB(u|D_2), \quad (\text{A.7})$$

then we will choose D_1 over D_2 . Likewise, if under D_1 we stand to gain more in the still probable worst case scenario than we stand to lose under the still probable best case scenario, that is, (A.6):

$$LB(u|D_1) - LB(u|D_2) > UB(u|D_2) - UB(u|D_1), \quad (\text{A.8})$$

then again we will choose D_1 over D_2 .

Note that the gains and losses in this discussion pertain to gains and losses on the utility dimension, not on the monetary outcome dimension. On the utility dimension the phenomenon of loss aversion, that is, the phenomenon that monetary losses may weigh heavier than equal monetary gains, has already been accounted for. Stated differently, the utility scale is a linear loss-aversion corrected scale for the moral value of monies.

Now, if we look at the scenarios (A.5) and (A.6), and the corresponding postulated rational, because intuitive, criteria of choice (A.7) and (A.8), then we see that we will choose D_1 over D_2 whenever we have that

$$LB(u|D_1) + UB(u|D_1) > LB(u|D_2) + UB(u|D_2). \quad (\text{A.9})$$

Moreover, this single criterion of choice is also consistent with the choosing of D_1 over D_2 in the scenarios (A.2), (A.3), and (A.4).

Note that if the decision inequality (A.9) goes to an equality:

$$LB(u|D_1) + UB(u|D_1) = LB(u|D_2) + UB(u|D_2). \quad (\text{A.10})$$

Then we have that we will be undecided when it comes to the decisions D_1 and D_2 .

Now, the k -sigma bounds in (A.1) translate to

$$[E(u|D_i) - k \text{std}(u|D_i), E(u|D_i) + k \text{std}(u|D_i)], \quad (\text{A.11})$$

for $i = 1, 2$, which, if substituted in (A.9), give the inequality

$$2E(u|D_1) > 2E(u|D_2), \quad (\text{A.12})$$

which brings us right back to Bernoulli's expected utility theory, as proposed in 1738, in which it is proposed that we choose that decision which maximizes the expectation value of the utility probability distributions [1].

Nonetheless, the criterion of choice, that the sum of the upper and lower bound should be maximized, as proposed here, will deviate from Bernoulli's initial 1738 proposal when the k -sigma intervals overshoot either the minimal or the maximal value of the utility probability distributions, or both.

Let a and b , respectively, be the minimal and maximal values of a given utility probability distribution. Then we may identify three additional cases, relative to (A.12):

$$LB^*(u) + UB^*(u) = \begin{cases} 2 E(u), & LB(u) \geq a, \quad UB(u) \leq b \\ a + UB(u), & LB(u) < a, \quad UB(u) \leq b \\ LB(u) + b, & LB(u) \geq a, \quad UB(u) > b \\ a + b, & LB(u) < a, \quad UB(u) > b \end{cases} \quad (\text{A.13})$$

These additional cases correspond, respectively, to scenarios where the k -sigma intervals (A.11) either undershoot, or overshoot, or both undershoot and overshoot the minimal and maximal values of a given utility probability distribution.

In the case of a lower confidence bound undershoot (e.g cases two and four) a too pessimistic worst case scenario is in play, and in the case of an upper confidence bound overshoot a too optimistic a best case scenario is in play, which is why we have to readjust these confidence bounds by replacing them with more realistic worst and best case scenarios (e.g the minimal and maximal values, a and b , of a given utility probability distribution).

It may be readily seen that any scalar multiple of (A.13) will retain the transitive ordering of the criterion of choice (A.13). Now, if we take as our scalar multiple $c = 1/2$, then our criterion of choice may be interpreted as a location measure of utility probability distribution:

$$\frac{LB^*(u) + UB^*(u)}{2} = \begin{cases} E(u), & LB(u) \geq a, \quad UB(u) \leq b \\ \frac{a+UB(u)}{2}, & LB(u) < a, \quad UB(u) \leq b \\ \frac{LB(u)+b}{2}, & LB(u) \geq a, \quad UB(u) > b \\ \frac{a+b}{2}, & LB(u) < a, \quad UB(u) > b \end{cases} \quad (\text{A.14})$$

where we note that the first case of (A.14) corresponds with Bernoulli's expected utility theory criterion of choice, whereas the fourth case corresponds with Hurwitz's criterion of choice with a balanced pessimism coefficient of $\alpha = 1/2$.

Note also that we have an added degree of freedom in our decision theory in that we may put an explicit premium on either caution or opportunity. For example, in

the absence of lower bound undershoot and upper bound overshoot, we have as the lower and upper bounds whose sum is to be maximized:

$$LB^*(u) = E(u) - k_1 \text{std}(u), \quad (\text{A.15})$$

and

$$UB^*(u) = E(u) + k_2 \text{std}(u). \quad (\text{A.16})$$

Then (A.15) and (A.16) sum to:

$$LB^*(u) + UB^*(u) = 2 E(u) + (k_2 - k_1) \text{std}(u). \quad (\text{A.17})$$

If in (A.17) we let $k_1 > k_2$, then we put a premium caution; alternatively, if we set $k_2 > k_1$, then we put a premium on opportunity; and if we let $k_1 = k_2$, then we have an equal trade-off between caution and opportunity taking, (A.13).